

# CS422 - Programming Language Design

## Lecture 11: Typed Languages - Type Inference

Grigore Roşu

Department of Computer Science  
University of Illinois at Urbana-Champaign

Checking that operations are applied on arguments of correct types statically has, as we have already discussed, **two major benefits**:

- Allows efficient implementations of programming languages by assuming untyped memory models and therefore removing the need for runtime consistency checks, and
- Gives rapid feedback to users, so they can correct errors at early stages.

These advantages are typically considered so major in the programming language community, that the *potential drawbacks* of static typing, namely

- Limiting the way programs are written by rejecting programs which do not type-check, and
- Adding typing information may be space and time consuming, are often ignored.

The first drawback is usually addressed by rewriting the code in a way that passes the type checking procedure. For example, programs passing a function as an argument to itself, which do not type check because of the recursive nature of the type of that function, can be replaced by equivalent programs using `let rec`.

It is, of course, desirable to require the user to provide *as little type information as possible* as part of a program. In general, depending on the application and domain of interest, there are quite *subtle trade-offs* between the amount of user-provided annotations and the degree of automation of the static analyzer.

In the previous lecture we discussed a simple static type checker which requires the user to provide a type for each declared name. Since some names can be functions taking functions as arguments, the amount and shape of these type annotations can be quite heavy. In some cases they may even exceed the size of the program itself. In this lecture we address the problem of reducing the amount of

required type information by devising automatic procedures that *infer the intended types from how the names are used*.

How much type information can be automatically and efficiently inferred depends again upon the particular domain of interest and its associated type system. In this lecture we will see an extreme fortunate situation, in which *all* the needed type information can be inferred automatically.

More precisely, we will discuss a classical procedure that *infers all the types* of all the declared names in our functional programming language. A similar technique is implemented as part of the **ML** and **OCaml** languages. By “type”, we here mean the language specific types, that is, those that were explicitly provided by users to the static type checker that we discussed.

## Type Inference

Programs written in the functional language discussed so far contain all the information that one needs in order to infer the intended type of each name. By carefully collecting and using all this information, one can type check programs without the need for the user to provide any auxiliary type information.

Suppose that one writes the expression  $x + y$  as part of a larger expression. Then one can infer from here that  $x$  and  $y$  are both intended to be integers, because the arithmetic operation  $+$  is defined only on integers! Similarly, if

```
if x then y else z
```

occurs in an expression, then one can deduce, thanks to the typing policy associated to conditionals, that  $x$  has type `bool` and that  $y$  and  $z$  have the same type. Moreover, from

```
if x then y else z + t
```

one can deduce that `z` and `t` are both of type `integer`, implying that `y` is also `integer`. Type inference and type checking work smoothly together, in the sense that one implicitly checks the types while inferring information. For example, if

```
if x then y else z + x
```

is seen then one knows that there is a typing error because the type of `x` *cannot* be both `integer` and `bool`.

The type of functions can also be deduced from the way the functions are used. For example, if one uses `f x y` in some program then one can infer that `f` takes two arguments. Moreover, if `f x x` is seen then one can additionally infer that `f`'s arguments have the same type.

There can be possible that the result type of a function as well as the types of its arguments can all be inferred from the context, without even analyzing the definition of the function. For example,

if one writes

```
(f x y) + x + let x = y in x
```

then one first infers that `f` must return an `int` because it is used as an argument of `+_`, then that `x` is an integer for the same reason, so the type of `f`'s first argument is `int`, and finally, because of the `let` which is used in a context of an `int`, that the type of `y` is also `int`. Therefore, the type of `f` must be `int -> int -> int`. In fact, the types of all the names occurring in the expression above were deduced by just analyzing carefully how they are used.

Let us now consider an entire expression which evaluates properly, for example one defining and using a factorial function:

```
let rec f n = if n eq 0
               then 1
               else n * f(n - 1)
in f 5
```

How can one automatically infer that this expression is type safe? Since `eq` takes two `int`s and returns a `bool`, one can deduce that the parameter of the function is `int`. Further, since the `f` occurs in the context of an integer and takes an expression which is supposed to be an integer as argument, `n - 1`, one can deduce that the type of that `f` is `int -> int`. Because of the semantics of `let rec`, the bound `f` will have the same type, so the type of the entire expression will be `int`.

One can infer/check the types differently, obtaining the same result.



## Polymorphic Functions

Let us consider a “projection” function which takes two arguments and always returns its first argument, such as `fun x y -> x`.

Without any additional type information, the best one can say about the type of this expression is that it is  $t_1 \rightarrow t_2 \rightarrow t_1$ , for some types  $t_1$  and  $t_2$ . This function can be now used in any context in which its result type equals that of its first argument, such as, `(int -> int) -> int -> (int -> int)`. Such functions are called *polymorphic*. Their type should be thought of as the *most general*, in the sense of the *least constrained*, type that one can associate them.

Let us consider several other polymorphic functions. For example, the type of the function `fun x y -> x y`, which applies its first argument, expected therefore to be a function, to its second argument, is  $(t_1 \rightarrow t_2) \rightarrow t_1 \rightarrow t_2$ .

How can one infer the most general type of an expression then, by just analyzing it syntactically? The process of doing this is called *type inference* and consists of two steps:

1. Collect information under the form of type parametric constraints by recursively analyzing the expression;
2. Solve those constraints.

## Collecting Type Information

In order to collect type information from an expression, we traverse the expression recursively, assign generic types to certain names and expressions, and then constrain those types. Let us consider, for example., the expression

```
fun x y -> if x eq x + 1 then x else y
```

and let us manually simulate the type information collecting

algorithm.

We have a function expression. Without additional information, the best we can do is to assume some generic types for its parameters and then calculate the type of its body. Let  $t_x$  and  $t_y$  be the generic types of  $x$  and  $y$ , respectively. If  $t_e$  is the type of function's body expression, then the type of the function will be

$$t_x \rightarrow t_y \rightarrow t_e.$$

Let us now calculate the type  $t_e$  while gathering type information as we traverse the body of the function.

The body of the function is a conditional, so we can now state a first series of constraints by analyzing the typing policy of the conditional: its condition is of type `bool` and its two branching expressions must have the same type, which will replace  $t_e$ .

Assuming that  $t_b$ ,  $t_1$  and  $t_2$  are the types of its condition and branching expressions, respectively, then we can state the type

constrains

$$t_b = \text{bool}, \text{ and}$$

$$t_1 = t_2.$$

We next calculate the types  $t_b$ ,  $t_1$  and  $t_2$ , collecting also the corresponding type information. The types  $t_1$  and  $t_2$  are easy to calculate because they can be simply extracted from the type environment:  $t_x$  and  $t_y$ , respectively (thus, the type equation  $t_1 = t_2$  above is in fact  $t_x = t_y$ ).

To type the condition  $x \text{ eq } x + 1$ , one needs to use the typing policy of `_eq_`: takes two arguments of type `int` and returns type `bool`. Thus  $t_b$  is the type `bool`, but two new type constraints are generated, namely

$$t_3 = \text{int}, \text{ and}$$

$$t_4 = \text{int},$$

where  $t_3$  and  $t_4$  are the types of the two subexpressions of `_eq_`.  $t_3$

evaluates to  $t_x$ ; in order to evaluate  $t_4$  one needs to apply the typing policy of  $_{+}$ : takes two `int` arguments and returns an `int`. These generate the constraint

$$t_x = \text{int}.$$

Therefore, after “walking” through all the expression and analyzing how names were used, we collected the following useful typing information:

$$t_x = t_y,$$

$$t_x = \text{int}, \text{ and}$$

the type of the original expression is  $t_x \rightarrow t_y \rightarrow t_x$ .

After we learn how to solve such type constraint equational systems we will be able to infer from here that the expression is correctly typed and its type is `int -> int -> int`.

Let us now consider the polymorphic function

```
fun x y -> (x y) + 1.
```

After assigning generic types  $t_x$  and  $t_y$  to its parameters, while applying the typing policy on `+_` one can infer that the type of the expression `(x y)`, say  $t_{(xy)}$ , must be `int`. Also, the result type of the function must be `int`. When calculating  $t_{(xy)}$ , by applying the typing policy for function application, which say that the type of `(F E)` for expressions `F` and `E` is `T` whenever the type of `F` is `Tp -> T` and that of `E` is `Tp`, one infers the type equational constraint  $t_x = t_y -> t_{(xy)}$ . We have accumulated the following type information:

$$t_{(xy)} = \text{int},$$

$$t_x = t_y -> t_{(xy)}, \text{ and}$$

$$\text{the type of the function is } t_x -> t_y -> \text{int}.$$

We will be able to infer from here that the type of the function is  $(t -> \text{int}) -> t -> \text{int}$ , so the function is *polymorphic in  $t$* .

## The Constraint Collecting Technique

One can collect type constraints in different ways. We will do it by recursively traversing the expression to type only once, storing a type environment as well as type constraints. If the expression is

- a *name* then we return its type from the type environment and collect no auxiliary type constraint;
- an integer or a bool constant, then we return its type and collect no constraint;
- an *arithmetic operator* then we recursively calculate the types of its operand expressions accumulating the corresponding constraints, then add the type constraints that their type is *int*, and then return *int*;
- a *boolean operator* then we act like before, but return *bool*;

- a *conditional* then calculate the types of its three subexpressions accumulating all the type constraints, then add two more constraints: one stating that the type of its first subexpression is `bool`, and another stating that the types of its other two subexpressions are equal.
- a *function* then generate a fresh generic type for its parameter, if any, bind it in the type environment accordingly, calculate the type of the body expression in the new environment accumulating all the constraints, and then return the corresponding function type;
- a *function application* then generate a fresh type  $t$ , calculate the type of the function expression, say  $t_f$ , and that of the argument expression, say  $t_p$  accumulating all the type constraints, and then add the type constraint  $t_f = t_p \rightarrow t$ .



- a *declaration* of names in a `let`, then calculate the types of the bound expressions, bind them to the corresponding names in the type environment, and then calculate and return the type of `let`'s body; there are no additional type constraints to accumulate (besides those in the bound expressions and body);
- a *declaration* of names in a `let rec`, then assign some generic types to these names in the type environment, calculate the types of bound expressions and accumulate the constraint(s) that these must be equal to the generic types assigned to `let rec`'s names, and finally calculate and return the type of `let rec`'s body;

- a *list expression*, then calculate the types of all the elements in the list accumulating all the corresponding constraints, and then add the additional constraints that all these types must be equal (let **T** be that type) and finally return the type of the list expression as **list T**;
- a *list operation*, such as **car**, **cdr**, **cons** or **null?**, then calculate the types of the corresponding subexpressions accumulating the constraints and then adding new, straightforward type constraints; for example, in the case of **cons(E,E')**, add the constraint that the type of **E'** must be **list T**, where **T** is the type of **E**;

- an *assignment*  $X := E$ , then calculate the types of  $X$  and  $E$ , and then add the constraint that these two types must be equal.
- a *sequential composition*  $E ; E'$ , then calculate the types of  $E$  and  $E'$  accumulating all the type constraints, then return the type of  $E'$  as the type of the composition;
- a *block*, then return the type `none` if the block is  $\{\}$ , and the type of  $E$  if the block is  $\{E\}$ ;
- a *loop while*  $\text{Cond Body}$ , then calculate the types of  $\text{Cond}$  and  $\text{Body}$ , add the constraint that the type of  $\text{Cond}$  must be `bool`, and then return the type `none`.

## Defining a Type Constraint Collector

We will use the same programming language and analysis tool equational definitional technique that we used to give semantics to **FUN** and to define its static type checker. Since we will not need type declarations (all types will be automatically inferred, we can just import the syntax of the language unchanged from the executable semantics definition, which we called `fun-syntax.maude`.

The **STATE** (module) of our type constraint collector will need to contain three attributes: a *type environment*, a *set of type equations*, and a *counter* needed to generate fresh types.

The constraint collector is defined in the same style as the other tools defined so far, namely inductively over the structure of the

syntax. More precisely, we will define an operation

`preType : Exp TypeState -> [Pair]`

which calculates the generic type of an expression in a given state, as well as a new state containing all the accumulated constraints. Since experience showed us that it is more convenient to handle lists of expressions in block, we will actually define the `preType` operator directly on lists of expressions.

We could have defined two distinct operations instead, one returning the type and the other the constraints (which can be regarded as “side effects”), as we did in the definition of the executable semantics. However, `preType` will combine both those operators into only one, thus allowing us to design a more compact and more easily understandable specification.

We call the type returned by `preType` a *pre-type*, because in order to calculate the final type we need to solve the equational constraints. We will first focus on defining `preType` and then on solving the type constraints.

In what follows, we only discuss the interesting aspects of the type constraint collector. The details can all be found in the file `fun-type-inference-semantics.maude` on the website.

## Defining Types

Like for the type checker, we introduce a specification defining types. Because of the more complex nature of type inference, and especially because of its ability to infer types in the context of polymorphism, we will need more structure on types than before. We declare a subsort `TypeVar` of the `Type` and a *type variable constructor* “`op t : Nat -> TypeVar`”, which will be used to generate fresh types as well as to represent polymorphic types:

```
fmod TYPE is protecting NAT .
  sorts Type TypeVar .  subsort TypeVar < Type .
  ops int bool unit none : -> Type .
  op _->_ : Type Type -> Type [gather(e E)] .
  op list_ : Type -> Type [prec 0] .
  op t : Nat -> TypeVar .
endfm
```

## Defining Type Equations

An *type equation* is a commutative pair of types. To handle blocks of constraints, we allow equations of lists of types (defined the usual way in a module `TYPE-LIST`) and translate them into sets of corresponding equations. The set of equations is kept small by removing the useless ones:

...

```

sorts Equation Equations .
subsorts Equation < Equations .
op empty : -> Equations .
op _=_ : TypeList TypeList -> Equation [comm] .
op _,_ : Equations Equations -> Equations [assoc comm id: empty] .
eq Eqn, Eqn = Eqn .
eq (T = T) = empty .
ceq (T,Tl = T',Tl') = (T = T'), (Tl = Tl') if Tl /= nil .

```



## Defining the State Infrastructure

The state of the type inferencer will need to contain three attributes: a type environment binding names to generic types, a set of equational type constraints, and a counter giving the next “free” type variable. The state can be easily defined by extending the state of the type checker discussed in the previous lecture with the two additional attributes.

```
fmod TYPE-STATE is extending TYPE-ENVIRONMENT .   extending EQUATIONS .
  sorts TypeStateAttribute TypeState .
  subsort TypeStateAttribute < TypeState .
  op empty : -> TypeState .
  op -- : TypeState TypeState -> TypeState [assoc comm id: empty] .
  op tenv : TypeEnv -> TypeStateAttribute .
  op eqns : Equations -> TypeStateAttribute .
  op nextType : Nat -> TypeStateAttribute .
endfm
```

Several standard helping operations are also needed and defined in the usual module `HELPING-OPERATIONS`. Only two helping operations are worthwhile discussing, namely:

```

op freshTypeVar : Nat TypeState -> Pair .
eq freshTypeVar(0, S) = {nil,S} .
ceq freshTypeVar(s(#), nextType(N) S) = {(t(N),T1), S1}
    if {T1,S1} := freshTypeVar(#, nextType(N + 1) S) .
...
op _&_ : TypeState Equations -> TypeState .
eq (eqns(Eqns) S) & Eqns' = eqns(Eqns, Eqns') S .

```

`freshTypeVar` takes a natural number `#` and a state and generates `#` “fresh” types together with the new state (containing an updated counter for fresh types). This operation will be intensively used in the sequel. `S & Eqns` collects the the equations in `Eqns` into the state `S`.

## Pre-Typing Generic Expressions

We next define the operation `preType` taking an expression and a state and returning a pair type-state, where the returned type is a temporary type calculated for the expression; that type will be “evaluated” into a concrete type only after solving the type constraints that are collected in the returned state.

The `preType` of integers and names is straightforward, because there is no constraint that needs to be collected:

```
eq preType(I, S) = {int, S} .
eq preType(X, S) = {S[X], S} .
```

As usual, we prefer to work with lists of expressions. Constraints must be properly accumulated when pre-typing lists of expressions:

```
ceq preType((E,E1), S) = {(T,T1), S1}
  if E1 /= nil /\ {T,S'} := preType(E,S) /\ {T1,S1} := preType(E1,S') .
```

## Pre-Typing Arithmetic and Boolean Expressions

When we pre-type an arithmetic or boolean operator, there are two things that we should consider, both derived from the typing policy of the language:

1. What are the expected types of its arguments and of its result;
2. How to propagate the type constraints.

...

```
ceq preType(E + E', S) = {int, S' & (T,T' = int,int)}
  if {(T,T'), S'} := preType((E,E'), S) .
```

...

```
ceq preType(E eq E', S) = {bool, S' & (T,T' = int,int)}
  if {(T,T'), S'} := preType((E,E'), S) .
```

...

## Pre-Typing the Conditional

The following definition is self-explanatory: pre-type the condition to **Tb** and the two branches to **T** and **T'**, respectively, and then add the appropriate type constraints:

```
ceq preType(if BE then E else E', S) = {T, S' & (Tb,T = bool,T')}
    if {(Tb,T,T'), S'} := preType((BE,E,E'), S) .
```

Those of you who have seen formal type systems before, most likely have seen a different notation, with a horizontal line separating the hypothesis (top) from conclusion (bottom). With that notation, the typing policy above would translate into something like:

$$\frac{S \vdash (BE, E, E') : \langle (Tb, T, T'), S' \rangle}{S \vdash (\text{if } BE \text{ then } E \text{ else } E') : \langle T, S' \& (Tb, T = \text{bool}, T') \rangle}$$

Books using the horizontal line notation to introduce type systems, typically contain (hard) exercises asking one to *implement* the

presented type system (in some programming language). In this class, we prefer the equational notation. It is not only equally compact and readable, but, more importantly, it is directly *executable*, so one does not need to implement anything.

Executability of specifications is a crucial advantage in designing a programming language, because, at design stage, one wants to focus on the major concepts/features introduced in the language and their relationships, rather than on low-level implementation details.

One of our major claims is that once one defines an executable semantics of a language or analysis tool using the advocated equational style, then one can easily *implement* that language or analysis tool in any programming language, if one really needs to. But an executable equational engine, such as Maude, can be so fast that one may not even be motivated to get a more efficient implementation (unless one's language becomes very popular and thus there is a public demand for a *very efficient* implementation).

## Pre-Typing Functions

As usual, we have to consider the two language constructs related to functions, namely function declaration and function invocation. For function declarations, the returned pre-type is the function type calculated as follows:

```
ceq preType(fun X -> E, S) = {Tp -> Tr, Sr[tenv <- tenv(S)]}
  if {Tp,Sp} := freshTypeVar(1,S)
  /\ {Tr, Sr} := preType(E, Sp[X <- Tp]) .
ceq preType(fun () -> E, S) = {unit -> Tr, Sr}
  if {Tr,Sr} := preType(E, S) .
```

Therefore, a “fresh” type is generated for the function’s parameter, if any, then the body of the function is pre-typed in the properly modified state, and then the resulting function type is returned. To be consistent with the semantics of **FUN**, the returned state *forgets* the parameter bindings. However, note that the type constraints

are not forgotten! They will be needed at the end of pre-typing.

Function invocations are a bit tricky, because fresh types need to be generated. More precisely, a fresh type is generated for each function invocation, and it is assumed to be the type of the result. Knowing that type, we can generate the appropriate constraints:

```
ceq preType(F E, S) = {Tr, Sr & (Tf = Tp -> Tr)}
  if {(Tf,Tp), Sf} := preType((F,E), S)
  /\ {Tr,Sr} := freshTypeVar(1,Sf) .
```

It is important to understand the need for the fresh type generated for the result of the function. One may argue that it is not needed, because one can simply pre-type  $F$ , which should pre-type to a function type  $Tp \rightarrow Tr$ , then  $E$  to some  $Tp'$ , and then return the type  $Tr$  and generate the obvious constraint  $Tp = Tp'$ . However, the problem is that  $E$  *may not pre-type to a function type*!

Consider, for example, the expression `fun x y -> (x y)`. When



pre-typing  $(x\ y)$ , the only thing known about  $x$  is that it has some generic type, say  $t(0)$ .  $y$  has also a generic type, say  $t(1)$ . Then how about the type of  $(x\ y)$ ? By generating a fresh type, say  $t(2)$ , for the result of the function application  $(x\ y)$ , we can conclude that the type of  $x$  should be  $t(1) \rightarrow t(2)$ , so we can generate the appropriate type constraint.

## Pre-Typing Let

In order to pre-type **let**, we first pre-type each bound expression and collect all the generated type constraints, then we assign fresh generic types to the names to be bound and add new corresponding type constraints, and then finally pre-type the body of the **let**.

```
ceq preType(let (X1,E1) in E, S) = {Te, Se[tenv <- tenv(S)]}
  if {Tel,Se1} := preType(E1,S)
  /\ {Te,Se} := preType(E, Se1[X1 <- Tel]) .
```

Notice that there are no additional type constraints to be propagated, and that, like for functions, the state which is returned after pre-typing **let** *forgets* the bindings added to pre-type its body, but retains all the generated type constraints as well as the new value of the fresh type counter!

## Pre-Typing Lists

As mentioned earlier, the typing policy of a list is that all its elements have the same type. In ensure this, we need to pre-type all the expressions in a list and to impose the constraints that all their types are equal. This can be done with one additional operator and a corresponding conditional equation:

```

op preType : Exp::List TypeState -> [Pair] .
ceq preType(E :: E::l, S) = {T, S'' & (T = T')}
  if {T,S'} := preType(E, S) /\ {T',S''} := preType(E::l, S') .

```

The returned type can be any of the types that occurred in the list; we chose the first one. Now a list can be pre-typed as follows:

```

ceq preType([E::l], S) = {list T, S'} if {T,S'} := preType(E::l, S) .

```

We have seen that the empty list should have a polymorphic type.

We can ensure that easily by generating a “fresh” type:

```
ceq preType([], S) = {list T?, S?} if {T?,S?} := freshTypeVar(1,S) .
```

The list operators can be pre-typed in a similar way. For example:

```
ceq preType(car(E), S) = {T?, S' & (list T? = T)}  
  if {T?,S?} := freshTypeVar(1,S) /\ {T,S'} := preType(E, S?) .
```

## Solving the Type Constraints

We can pre-type all the other language constructs in a very similar way (this will be part of your homework). Thus, the operator `preType` will take an expression and return a pair  $\{T, S\}$ , where  $T$  is a term of sort `Type` and `eqns(S)` contains all the *type constraints* accumulated by traversing the program and applying the typing policy. We refer to terms  $t(0)$ ,  $t(1)$ , ..., as *type variables*.

We should understand the result  $\{T, S\}$  of `preType(E)` as follows:

The (final) type of  $E$  will be  $T$  in which all the type variables will be *substituted* by the types obtained after *solving* the system of equations in  $S$ .  $E$  is *not guaranteed* a type *a priori*! If any *conflict* is detected while solving the system then we say that the process of typing  $E$  failed, and we conclude that the expression is not correctly typed.

If **Eqns** is a set of type equations, that is, a term of sort **Equations**, and **T** is a type potentially involving type variables, then we let **Eqns [T]** denote the type obtained after solving **Eqns** and then substituting the type variables accordingly in **T**.

Once such a magic operation **\_[\_] : Equations Type -> [Type]** is defined, we can easily find the desired type of an expression:

**ceq type(E) = eqns(S) [T] if {T,S} := preType(E) .**

We next describe a simple equational technique to define the operation **\_[\_] : Equations Type -> [Type]**.

**Homework Exercise 1** *Complete the definition of the type inferencer in **fun-type-inference-semantics.maude**. You need to pre-type the remaining language constructs and then to formalize the technique that will be next presented.*

# The Type Inference Procedure by Examples

## Example 1

Let us consider the following expression:

```
(fun x y -> x y) (fun z -> 2 * z * z * 3) 3
```

After pre-typing (do it, using `preType` on it), we get the pre-type `t(5)` and a state containing the four equations:

```
int = t(3),  
t(0) = t(1) -> t(2),  
t(4) = int -> t(5),  
t(0) -> t(1) -> t(2) = (t(3) -> int) -> t(4)
```

How can we solve the system of equations above? We can immediately notice that `t(3) = int`, so we can just substitute `t(3)` accordingly in the other equations, obtaining:

$$t(0) = t(1) \rightarrow t(2),$$

$$t(4) = \text{int} \rightarrow t(5),$$

$$t(0) \rightarrow t(1) \rightarrow t(2) = (\text{int} \rightarrow \text{int}) \rightarrow t(4)$$

Moreover, since  $t(0) = t(1) \rightarrow t(2)$ , by substitution we obtain

$$t(4) = \text{int} \rightarrow t(5),$$

$$(t(1) \rightarrow t(2)) \rightarrow t(1) \rightarrow t(2) = (\text{int} \rightarrow \text{int}) \rightarrow t(4)$$

And further, by substituting  $t(4)$  into the last equation we get

$$(t(1) \rightarrow t(2)) \rightarrow t(1) \rightarrow t(2) = (\text{int} \rightarrow \text{int}) \rightarrow \text{int} \rightarrow t(5)$$

Both types in the equation above are function types, so their source and target domains must be equal. This observation allows us to generate two other type equations, namely:

$$(t(1) \rightarrow t(2)) \rightarrow t(1) = (\text{int} \rightarrow \text{int}) \rightarrow \text{int},$$

$$t(2) = t(5).$$

Now, substituting  $t(2)$  by  $t(5)$  in the first equation we get:

$$(t(1) \rightarrow t(5)) \rightarrow t(1) = (\text{int} \rightarrow \text{int}) \rightarrow \text{int}$$



that is, an equality of function types, which yields the following:

$$\begin{aligned} t(1) \rightarrow t(5) &= \text{int} \rightarrow \text{int}, \\ t(1) &= \text{int}. \end{aligned}$$

Substituting  $t(1)$  in the first equation and then equating the source and target types of the function type, we get

$$\begin{aligned} \text{int} &= \text{int}, \\ t(5) &= \text{int}. \end{aligned}$$

The first equation is useless and will be promptly removed by the simplifying rule  $(T = T) = \text{none}$  in **EQUATIONS**. The second equation gives the desired type of our expression,  $\text{int}$ .

## Example 2

Let us now see an example expression which cannot be typed, e.g.:

```
let x = fun x -> x in (x x)
```

After pre-typing, we get the type  $t(1)$  constrained by:

$$t(0) \rightarrow t(0) = (t(0) \rightarrow t(0)) \rightarrow t(1)$$

We can already see that this equation is problematic, but let us just follow the procedure blindly to see what happens. Since both types in the equation are function types, their source and target types must be equal, so we can generate the following two equations:

$$t(0) = t(0) \rightarrow t(0),$$

$$t(0) = t(1)$$

The first equation is obviously problematic because of its “recursive” nature, but let us ignore it and substitute the second into the first, obtaining:

$t(1) = t(1) \rightarrow t(1).$

There is no way to continue. The variable  $t(1)$  cannot be assigned a type, so the original expression *cannot be typed*.

## Unification

The technique that we used to solve the type equations is called *unification*. In general the equational unification problem can be stated as follows, where we only consider the mono-sorted case. Our particular unification problem fits this framework, because we can consider that we have only one sort, **Type**.

Let  $\Sigma$  be a signature over only one sort, i.e., a set of operations like in **Maude**, let  $X$  be a finite set of variables, and let

$$\mathcal{E} = \{t_1 = t'_1, \dots, t_n = t'_n \mid t_1, t'_1, \dots, t_n, t'_n \in T_\Sigma(X)\}$$

be a finite set of pairs of  $\Sigma$ -terms over variables in  $X$ , which from now on we call *equations*. Notice, however, that these are different from the standard equations in equational logic, which are quantified universally.

A map, or a substitution,  $\theta : X \rightarrow T_\Sigma(X)$  is called a *unifier* of  $\mathcal{E}$  if and only if  $\theta^*(t_i) = \theta^*(t'_i)$  for all  $1 \leq i \leq n$ , where  $\theta^* : T_\Sigma(X) \rightarrow T_\Sigma(X)$  is the natural extension of  $\theta : X \rightarrow T_\Sigma(X)$  to entire terms, by substituting each variable  $x$  by its corresponding term,  $\theta(x)$ .

A unifier  $\theta : X \rightarrow T_\Sigma(X)$  is *more general* than  $\varphi : X \rightarrow T_\Sigma(X)$  when  $\varphi$  can be obtained from  $\theta$  by further substitutions; formally, when there is some other map, say  $\rho : X \rightarrow T_\Sigma(X)$ , such that  $\varphi = \theta; \rho^*$ , where the composition of function was written sequentially. Let us consider again our special case signature, that of types, and the equations  $\mathcal{E}$ :

$$t(0) = t(1) \rightarrow t(2),$$

$$t(0) \rightarrow t(1) \rightarrow t(2) = (t(3) \rightarrow t(3)) \rightarrow (t(4) \rightarrow t(4)) \rightarrow t(5).$$

Here  $t(0)$ , ...,  $t(5)$  are seen as variables in  $X$ .

One unifier for these equations, say  $\varphi$ , takes  $t(0)$  to

$$(int \rightarrow int) \rightarrow (int \rightarrow int),$$

$t(1)$ ,  $t(2)$  and  $t(3)$  to  $int \rightarrow int$ , and  $t(4)$  to  $int$ . Another unifier, say  $\theta$ , takes  $t(0)$  to

$$(t(4) \rightarrow t(4)) \rightarrow (t(4) \rightarrow t(4)),$$

and  $t(1)$ ,  $t(2)$  and  $t(3)$  to  $t(4) \rightarrow t(4)$ . Then it is clear that  $\theta$  is *more general* than  $\varphi$ , because  $\varphi = \theta; \rho^*$ , where  $\rho$  simply takes  $t(4)$  to  $int$ .

If a set of equations  $\mathcal{E}$  admits a unifier, they are called *unifiable*. Note that there can be situations, like the ones we have already seen for our special case of signature of types, when a set of equations is not unifiable.

## Finding the Most General Unifier

When a set of equations  $\mathcal{E}$  is unifiable, its *most general unifier*, written  $mgu(\mathcal{E})$  is one which is more general than any other unifier of  $\mathcal{E}$ . Note that typically there are more than one  $mgu$ .

We will next discuss a general procedure for finding an  $mgu$  for a set of equations  $\mathcal{E}$  over an arbitrary mono-sorted signature  $\Sigma$ . Since we actually need to apply that  $mgu$  to the type calculated by the `preType` operation for the expression to type, we will provide a technique that directly calculates  $\mathcal{E}[t]$  for any term  $t$ , which calculates the term after applying the substitution  $mgu(\mathcal{E})$  to  $t$ , that is,  $(mgu(\mathcal{E}))^*(t)$ .

The technique is in fact quite simple. It can be defined entirely equationally, but in order to make the desired direction clear we write them as rewriting rules (going from left to right). It consists

of iteratively applying the following two steps to  $\mathcal{E}[t]$ :

1.  $(x = u, \mathcal{E})[t] \rightarrow \text{subst}(x, u, \mathcal{E})[\text{subst}(x, u, t)]$ , if  $x$  is some variable in  $X$ ,  $u$  is a term in  $T_\Sigma(X)$  which does *not* contain any occurrence of  $x$ , and  $\text{subst}(x, u, \mathcal{E})$  and  $\text{subst}(x, u, t)$  substitute each occurrence of  $x$  in  $\mathcal{E}$  and  $t$ , respectively, by  $u$ ;
2.  $(\sigma(t_1, \dots, t_k) = \sigma(t'_1, \dots, t'_k), \mathcal{E})[t] \rightarrow (t_1 = t'_1, \dots, t_k = t'_k, \mathcal{E})[t]$ , if  $\sigma \in \Sigma$  is some operation of  $k$  arguments.

**Theorem.** *When none of the steps above can be applied anymore, we obtain a potentially modified final set of equations and a final term, say  $\mathcal{E}_f[t_f]$ . If  $\mathcal{E}_f$  is empty then  $t_f$  is  $(\text{mgu}(\mathcal{E}))^*(t)$ , so it can be returned as the type of the original expression. If  $\mathcal{E}_f$  is not empty then  $\mathcal{E}$  is not unifiable, so the original expression cannot be typed.*



## Let Polymorphism

Unfortunately, our current polymorphic type inference technique is not as general as it could be. Suppose, for example, that one wants to define a function with the purpose of using it on expressions of various types, such as an identity function. Since the identity function types to a polymorphic type,  $t(0) \rightarrow t(0)$  or some similar type, one would naturally want to use it in different contexts. After all, that's the purpose of polymorphism.

Unfortunately, our current type inferencer *cannot* type the following expression, even though it correctly evaluates to `int(1)`:

```
let f = fun x -> x
in if f true then f 1 else f 2
```

The reason is that `f true` will already enforce, via the type constraints, the type of `f` to be `true -> true`; then `f 1` and `f 2`

will enforce the type of  $f$  to be  $\text{int} \rightarrow \text{int}$ , thus leading to a contradiction. Indeed, our type inferencer will reduce  $\mathcal{E}[t]$ , where  $\mathcal{E}$  is the set of constraints and  $t$  is the pre-type of this expression, to  $\mathcal{E}_f[t_f]$ , where  $\mathcal{E}_f$  is the equation  $\text{bool}=\text{int}$  and  $t_f$  is  $\text{int}$ ; since  $\mathcal{E}_f$  is not empty, the original expression cannot be typed.

The example above may look artificial, because one does not have a need to define an identity function. Consider a length function instead, which is, naturally, intended to work on lists of any type:

```
fun x -> let c = 0
        in {
            while(not null?(x))
            {
                x := cdr(x) ;
                c := c + 1
            } ;
            c
        }
```

As expected, our type inferencer will type this expression to a polymorphic type, `list t(0) -> int`. However, if one uses it on a list of concrete type, then `t(0)` will be unified with that type, thus preventing one from using the length function on other lists of different types. For example,

```
let l = fun x -> let c = 0
                in {
                    while(not null?(x))
                    {
                        x := cdr(x) ;
                        c := c + 1
                    } ;
                    c
                }
in l [5 :: 3 :: 1] + 1 [true :: false :: true]
```

will not type, even though it correctly evaluates to `int(6)`.

This leads to the following natural question, that a programming language designer wishing a type system on top of his/her language must address: what should one do when valid programs are rejected by the type system? One unsatisfactory possibility is, of course, to require the programmers write their programs differently, so that they will be accepted by the type system. Another, better possibility is to improve the type system to accept a larger class of programs.

Since *reusing* is the main motivation for polymorphism, it would practically make no sense to disallow the use of polymorphic expressions in different contexts. Therefore, programming language scientists fixed the problem above by introducing the important notion of *let-polymorphism*.

The idea is to type the **let** language constructs differently. Currently, we first type all the binding expressions to some (pre-)types, then we bind those types to the corresponding names

into the type environment, and finally we type the body expression in the new environment. This technique is inherently problematic with respect to reuse, because a (polymorphic) binding expression is enforced to have two different types if used in two different contexts, which is not possible.

The idea of *let-polymorphism* is to first *substitute* all the free occurrences of the bound names into the body of the **let** expression by the corresponding binding expressions, and then to type directly the body expression. This way, each use of a bound name will be typed locally, thus solving our problem.

**Exercise 1** *Modify our current type inferencer to include let-polymorphism. How about **let rec**?*

However, a complexity problem is introduced this way: the size of the code can grow exponentially and the binding expressions will be (pre-)typed over and over again, even though each time they

will be typed to “almost” the same type (their types will differ by only the names of the type variables). There is a partial solution to this problem, too, but that is a more advanced topic that will be covered in CS522 next semester.

## Type Systems May Reject Correct Programs

While static type checkers/inferencers offer a great help to programmers to catch (trivial) errors at early stages in their programs and to programming language implementers to generate more efficient code (by assuming the untyped memory model), they have an inherent drawback: they *disallow correct programs which do not type check*. In the previous section we saw a fortunate case where a type system could be relatively easily extended (let-polymorphism), to allow polymorphic definitions to be used as desired. Unfortunately, due to undecidability reasons, for any given

(strong) type system, there will always be programs which can evaluate correctly but which will not type check.

In our context, for example, consider the following expression/program for calculating the length of a list *without* using `let rec` or loops:

```
let t f x = if null?(x) then 0 else 1 + f f cdr(x)
in let l x = t t x
    in 1 [4 :: 7 :: 2 :: 1 :: 9]
```

The idea in the program/expression above is to pass the function as an argument to itself; this way, it will be available, or “seen”, in its own body, thus simulating the behavior of `let rec` without using it explicitly. That is how (functional) programmers wrote recursive programs in the early days of functional programming, when the importance of static scoping was relatively accepted but no clean support for recursion was available. It, indeed, evaluates to `int(5)` but does not type. Try it!

One may wrongly think that the above does not type because our type system does not include let-polymorphism. To see that let-polymorphism is not the problem here, one can apply the substitutions by hand and thus eliminate the `let` constructs entirely:

```
(fun f x -> if null?(x) then 0 else 1 + f f cdr(x))
  (fun f x -> if null?(x) then 0 else 1 + f f cdr(x))
    [4 :: 7 :: 2 :: 1 :: 9]
```

The above will again evaluate to `int(5)`, but will not type. The set of constraints is reduced to `t(6) = t(6) -> list int -> int`, reflecting the fact that there is some expression whose type has a recursive behavior. Indeed, the type of `f` (any of them) cannot be inferred due to its recursive behavior.



In fact, the problematic expression is

```
fun f x -> if null?(x) then 0 else 1 + f f cdr(x)
```

which cannot be typed because of the circularity in the type of `f`.

No typed functional languages can actually type this expression. In `ML`, for example, the same expression, which in `ML` syntax is

```
fn (f,x) => if null(x) then 0 else 1 + f f tl(x)
```

generates an error message of the form:

```
t.sml:16.36-16.49 Error: operator is not a function [circularity]
operator: 'Z
in expression:
  f f
```